Decreased Time Delay and Security Enhancement Recommendations for AMI Smart Meter Networks[1]

Coalton Bennett and Steven B. Wicker
Department of Electrical and Computer Engineering, Cornell University
E-mail: cb322@cornell.edu

***Abstract*—In this paper we consider a variety of real world AMI smart meter networking scenarios. We examine the network performance as a function the following parameters: the size of the network, node scheduling, and polling interval (period) under normal conditions. Then we introduce malevolent agents to the network to demonstrate the effects of their actions. Although there are several inherent security related issues involving wireless network the most relevant and attacks to launch against the de facto networking protocols used by AMI smart meters is a black hole attack. Focusing on the black hole attack we demonstrate, through simulations, how to avoid these attacks by creating dedicated paths between the source (smart meter) and the sink collecting the information (access point). In addition to this, we simulate a network of nodes that use a hybrid routing protocol when time sensitive data is needed in order to ensure that outages do not occur due to excess demand.**

***Index Terms*—Advanced Metering Infrastructure (AMI), Zigbee, Ad Hoc On Demand Distance Vector (AODV), smart meter, Multiprotocol Label Switching (MPLS), demand response, Automatic Meter Reading (AMR), Meter Data Management Systems (MDMS)**

I. Nomenclature

A nomenclature list, if needed, should precede the Introduction.

II. Introduction

WITHIN the past few years the utility industry has undergone, and will continue to undergo, a rapid transformation. An array of new technologies for transmission and distribution systems have made their way into a number of electric utilities service territories.

The combination of Phasor Measurement Units and Global Positioning Systems time stamps have made it possible to measure the magnitude and phase angle of voltage and current waveforms at multiple locations in the transmission system at the same time. Demand response programs incorporating the use of smart meters to provide customers and utilities with real time power measurements have also been deployed in a number of pilot projects across America to help with customer saving and efficient management of utility resources. Even small: software, hardware, and system level development companies [1][2] have begun creating solutions for customers looking to reduce their energy bills. Google has developed a smart meter which uses the Noninstrusive Load Monitoring Algorithms developed at MIT, to provide customers with an in home display of the energy being consumed in real time. These companies along with several others are forerunners in an industry, which is projected to have the same if not more growth potential as the internet [3].

The majority of devices plugged into the grid generate a response based on one of two different stimuli. The first is a response generated by the electrical system without human intervention (e.g. switch opening within the distribution network as a result of the current or voltage being above a certain threshold). The second is a response generated from human intervention (e.g. adjusting the thermostat). Incorporating wireless sensor networks with legacy systems, infrastructure, and devices would allow these entities to make self-informed decisions using the two aforementioned stimuli. As an example consider a scenario in which each customer's meter and a subset of appliances have been upgraded to include this additional functionality. The additional functionality provided by the sensors within the home will enable appliances to make informed decisions about toggling the states of the devices based on price, time of day, or other parameters designated by the customer. The appliances that the customer decides to enroll in a demand response program will be connected to the Internet either directly through a personal computer or through their meter. In either case this effectively increases the number of endpoints connected to the Internet by a substantial amount. The question then becomes, how should all of this information be handled? Is it best to use the customer's home computer, use Broadband over Powerline, the dial up modems used for AMR, or a wireless solution?

The Zigbee Alliance has created an open standard for metering and appliance communication, which is largely based on the Zigbee Pro and 802.15.4 standards [4]. This is one of several excellent solutions that have been proposed to deal with the issue of sending customer data securely and reliably to the utilities meter data management systems (MDMS). Although there are a number of other communication platforms that could be used, some better than the Zigbee/802.15.4 platform, little to no effort has been put forth by vendors, utilities, or government agencies to research these other solutions. As a result Zigbee has been standardized as the solution of choice for utilities awarded stimulus money to upgrade their existing meters to be compliant with AMI system specification requirements.

Zigbee was originally developed for Low Rate Wireless Personal Area Networks. While it is projected that meters will only be polled every 15 minutes the control overhead associated with establishing a connection between the meter (source) and the destination (sink) is extremely high. This is

Coalton Bennett is a Ph.D candidate at Cornell University, Department of Electrical and Computer Engineering, Ithaca NY 14850 (e-mail: cb322@cornell.edu)

due to the fact that Zigbee uses both cluter tree and AODV routing protocols. Cluster tree routing is only used on Reduced Functional Devices (RFDs), which are usually battery operated or resource constrained devices. AODV was created for use by Full Functional Devices (FFDs), which are less resource constrained than RFDs because they are usually tethered to a power supply. Cluster tree routing suffers from tremendous delays due to the multiteered nature of the protocol, whereas AODV is much faster because routes are established as needed. Thus it is envisioned that the routing protocol will be pure AODV. Like any other routing protocol there are shortcomings or tradeoffs that must be considered. In the case of AODV it is the control overhead associated with the establishment of a connection between a meter and the utility. If vendors and utilities are not aware of this, which we are inclined to believe that they are not [5], then during times of increased demand the utilities demand response system will experience tremendous bottlenecking. To this end we suggest that a new layer be added between the Zigbee routing layer and the 802.15.4 Medium Access Control similar to the 2.5 Layer used in Multiprotocol Label Switching networks to decrease the end-to-end delays for data packets.

The concept of using labels to decrease end-to-end delays has been explored in several works [10][11] and the results show that the end-to-end delays decrease with use of a label switching protocol. The use of labels will enable the utilities network engineer to poll the meters and receive responses at sub 15 minute intervals without congesting the network during mission critical situations when demand begins to increase beyond a system sustainable limit.

III. Motivation

The introduction of smart meters has created a number of interesting networking problems. In our framework we view each meter as a sensor and actuator. Each meter measures the amount of real and reactive power and then adjusts these parameters based on user-defined values. In this sense the network can be viewed as a traditional sensor networking problem. Unlike traditional wireless networks, wireless sensor networks, usually operate under a variety constraints. These include, but are not limited to: bandwidth, battery power, limiting jitter (in time sensitive applications), and physical tampering. Usually the objective in a wireless sensor network is to maximize the network lifetime by developing routing protocols that minimize battery usage, and are robust in the face of: network, data link, and physical layer attacks. Although battery power is not an issue for this particular application, most of the other issues listed above are also concerns for a smart meter network as well. As a result perhaps the correct way to design and analyze the smart meter network is to view each meter as a node in a wireless sensor network. Wireless networks are similar in nature to wireless sensor networks, but some aspects of the two networks differ tremendously. For instance most nodes in a traditional wireless network (802.11 WLAN) can tolerate a certain amount of jitter and still provide a reasonable service to the user. This is usually due to the fact that the hardware is slightly more sophisticated and also the nodes communicate directly with the access points or routers in the network. With a network of nodes randomly dispersed over a, possibly non-uniform cross sectional area of several kilometers, using similar hardware and a direct line of sight connection might be costly given the options available (WiMAX) [6] but would decrease the surface area of attack. As a result a low rate wireless personal network option was chosen with a multihop routing scheme. The first major drawback of using a multihop routing protocol is the increase in surface area to attack. The additional nodes that a packet must travers makes it extremely easy for an adversary anywhere in the network between the source and the sink to launch a series of network layer attacks (black hole, grayhole, sybil attacks) [7]. Some vendors have taken security into consideration when designing their meters, by incorporating the AES-128,256 encryption standards. Whereas with other meters [1], it has been shown that it is possible to view neighbor energy consumption data thus providing a vein into the network from which a full scale attack could launched thus devastating entire portions of the network. In either case if proper precautions are not taken into consideration when addressing the inherent security issues at the network layer, then portions of the network are left vulnerable to attack. These issues serve as our motivation to develop a solution to the problem(s) that exist at the network layer for some of the meters being used by the nation's largest utilities.

IV. Problem Description

The Zigbee protocol uses a hybrid routing scheme, which incorporates elements of cluter-tree and AODV [8]. Zigbee devices are usually divided into two categories: Full Functional Devices and Reduced Functional Devices. FFDs are usually tethered devices and use both Cluster Tree (CT) routing and implement a variant of AODV called AODVjr, although the underlying principals are still the same. Whereas RFDs use CT routing only because it is less memory intensive and RFDs are usually battery powered. CT routing not only requires more memory, but the delays associated with it being to increase exponentially with the number of nodes using the routing protocol. As a result AODV should be the preferred choice for networks with tethered devices like the smart meters.

A. AODV

AODV is a very straight forward routing scheme, and was originally developed for Mobile Ad Hoc Networks where users in the network use intermediate nodes to route information to the desired destination. The AODV communication process and can be divided into three phases. The first phase is the discovery phase, whereby a Route Request (RREQ) is sent from the source to the desired destination. Each intermediate node will forward the RREQ if it does not posses a route to the destination that is recent enough as designated by a sequence number in the

RREQ packet.[1] If an intermediate node has a sequence number which is greater than the sequence number in the packet then the intermediate node will not forward the RREQ and further but instead will send a Route Reply (RREP) message back to the requesting (source) node along the same path from whence the RREQ came.[2] The second phase consists of updating the destination sequence numbers in the routing tables of the intermediate nodes for the destinations that they have received RREQs for. The third phase consists of sending the data once the route has been established between the source and the destination.

This process usually repeats itself if HELLO messages are not sent within a specified time period. HELLO messages are used by all nodes in the network to maintain connectivity with neighboring nodes, and to update routing table information for the destinations in their routing table. They are normally used in applications with nodes that demonstrate a high degree of random mobility and require a constant connection with a destination. Smart meters do not have the same demands and therefore do not require the usage of HELLO messages at regular intervals. In fact the discovery phase of the algorithm need only be performed once because the network is stationary. Furthermore repeated flooding of the network with RREQ packets has two potential drawbacks. The first is that the network will experience sever delays because a significant amount of network resources, in this case time, will be spent setting up and tearing down routes from the source to the destination. Second black hole attacks can easily be launched when repeated RREQs are sent.[3] This is perhaps the most common and easiest attack to launch, along with a destination sequence number attack, against a network that uses an on demand routing protocol. The easiest way to combat such an attack is to construct a dedicated path between the source and the destination. Otherwise a compromised node will begin to attract all data packets for which a node uses it as a route to the destination and they will be discarded. This could potentially have a cascading affect leading to network segmentation or even worse possible system wide outages. It is only necessary to use the discovery phase once and then begin sending data along the constructed path. Although a dedicated path does not protect against a node that will drop packets routed through it along the dedicated path, it will effectively reduce the nodes ability to *attract* more packets than it normally should because RREQs and RREPs will only be sent during the discovery phase.

[1] The destination sequence number is an indicator of how "fresh" a route is to the requested destination.

[2] This marks the establishment of a dedicated path from the source to the destination.

[3] Black hole attack-when a malicious or compromised node receives a RREQ it immediately sends a RREP back to the node issuing the RREQ with a destination sequence number greater than the one in the RREQ so that the issuing node sends it's data packets then the compromised node drops the data packets.

B. Zigbee/802.15.4

A network of 2500 Zigbee/802.15.4 radios, operating at a the maximum data rate of 250kbps, being polled every 15 minutes will only be able to transmit a file with at most 11,250 bytes which corresponds with 88 packets. For a file of this size each node will only be able to transmit at most 244 packets per second, which only gives the meter approximately 4 milliseconds to deliver a packet of this size. The size of the network was chosen to reflect an ongoing pilot project of equal size or greater. Given this value, and depending on the size of the network and the distance between homes, the amount of time available for retransmission by intermediate nodes is almost infinitesimal. If additional control messages are sent throughout the network then time spent in preparation for sending a packet as well as the time spent sending the packet could potentially create bottlenecks, requiring nodes to retransmit their data. The situation becomes even worse as the polling period increases requiring nodes to transmit the same or possibly even more data at a faster clip. In the next section we show how decreasing node processing time associated with packet forwarding increases the amount of time available for packet forwarding and retransmissions. This becomes absolutely essential for networks of the size described above with polling periods less than the de facto 15-minute interval.

V. Approach & Results

The network described in the previous section was simulated using OPNET [9]. We used the standard IEEE 802.11 radio, and made as many modifications as OPNET would allow resembling the IEEE 802.15.4 radio.

PHY

802.11	802.15.4
DSSS	DSSS
BPSK	BPSK
1 Mbps	250kbps

MAC

802.11	802.15.4
CSMA/CA (DCF)	CSMA/CA (Exponential Backoff)
RTS/CTS	N/A

The major differences between the two protocols happens to be the data rate, which is embedded in the source code and would require the creation of a new process model in order to create an exact replica of the 802.15.4 standard. Also 802.11 uses a Distributed Coordination Function in order to determine the backoff period, and differs only from the exponential backoff mechanism used by 802.15.4 in that the binary exponential random variable is multiplied by a slot time. Where the slot time is defined as a uniform period of time that elapses during which nodes can send messages (e.g. 2 millisecond per slot). We considered this to be an

unsubstantial amount of time to warrant any major differences between the two protocols. The RTS/CTS mechanism is disabled by default in OPNET enabling a pure CSMA/CA MAC protocol. The routing protocol selected was AODV without cluster tree routing. The polling period was changed to once every second to emphasize the effect that a black hole attack can have on network performance. Because the data rate for the 802.11 radios could not be changed we decided to transmit one file per node containing the largest number of bits that could be supported by a radio with a data rate of 1Mbps over a 60 minute interval. We used 128 bytes per packet[4], with a constant packet rate of one packet per second. Each node was randomly deployed over a uniform grid with an area of 1 square kilometer, and the access point (sink) was placed in the middle.

To simulate the black hole attack, we randomly selected a node in the network to launch the attack on, and changed the parameters of the nodes downstream from the selected node to ensure that all data packets were dropped. The AODV ACTIVE_ROUTE_TIMEOUT field was set to zero. An active route timeout of zero will force the nodes downstream from the selected node to send it's highest priority packets first. Packets arriving from the selected node will be buffered until the node has sent it's own packets. As a result all of the packets being buffered from the selected node will be dropped, if the downstream node has not processed the packets before ACTIVE_ROUTE_TIMEOUT period expires. In addition to this the processing time per packet was increased for each node downstream from the selected node. The packet queue size was also set to zero to ensure that any packet received, and queued, from the randomly selected node was automatically dropped.

The network is shown in Figure 1. The node under attack is in the lower right hand corner and the downstream nodes marked by red Xs collaboratively launch a black hole attack. After changing the attributes listed above, the throughput of the randomly selected node decreased to zero Figure 2. The node continues to send packets even though none of it's packets are being delivered. Clearly this indicates the severity of such an attack. This demonstration shows how a collaborative black hole attack can limit, if not completely reduce, the throughput to zero, when AODV control messages are used to maintain link connectivity.

This can be resolved though if a dedicated path is created during the network discovery phase.[5] This can be seen by the increase in throughput for the node originally under attack in Figure 3. As can be seen the packets actually reach the intended destination. The end-to-end delays can also be decreased as shown in Figures 4 and 5 if the system operators deem it necessary to poll customer premises at a faster clip. This is accomplished by using a modified MPLS layer 2.5 protocol [12]. Because our objective is to show that the addition of the layer 2.5 protocol will decrease end-to-end delays we simply increased the processing speed for packet forwarding. In essence this is exactly what the MPLS protocol does. At each node a label is removed and another one is added to the header in the packet. The nodes only read the label as opposed to the entire packet, and perform a table look up based on the label assignments. As a result this decreases the amount of time and resources spent on processing a packet in order to determine the next hop for the packet and subsequently decreases the end-to-end delays.[6] Figure 4 is the end-to-end delay for all IP traffic in the network from source to sink when a label switching protocol is not used. Figure 5 is the same statistic for a network using a label switching protocol. At the beginning of the simulation both networks exhibit a sharp increase in the traffic delay however with the label switching protocol the network begins to show signs of a decrease in delay. The network that is not using the label switching protocol experiences an increase in delay. Although the increase in delay is small, given the various fluctuations: in voltage, real and reactive power, as well as changing power factor values within a matter of seconds portions of a utilities service territory could be without power. Thus warranting the use of a protocol for time sensitive applications. When system operators need to issue warnings or send messages to customers indicating that they should reduce their demand a hybrid protocol should be used.

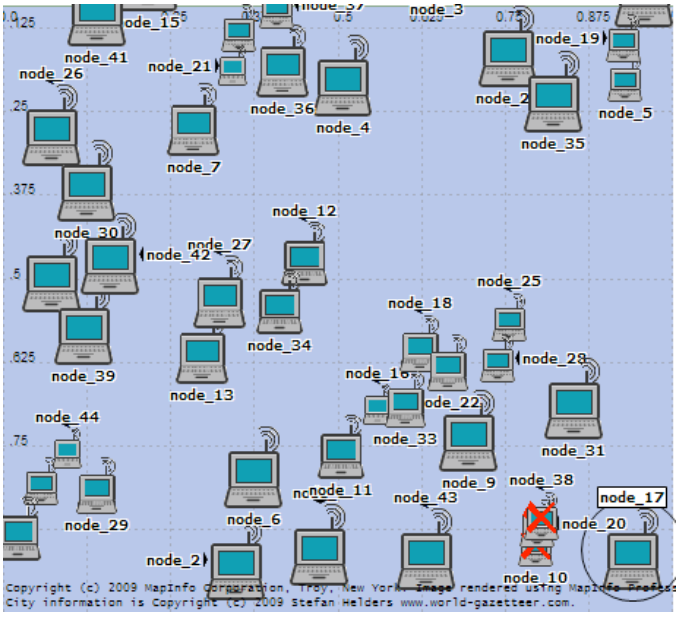

Figure 1.

[4] The maximum number of bytes per packet as specified by the Zigbee/802.15.4 standard

[5] Implicitly we assume that the black hole attack has been launched after the network has been established. This is a reasonable assumption, because suspicious behavior is more likely not to go unnoticed during the initial roll out.

[6] Readers are encouraged to read RFC 3031 for detailed information about the protocol

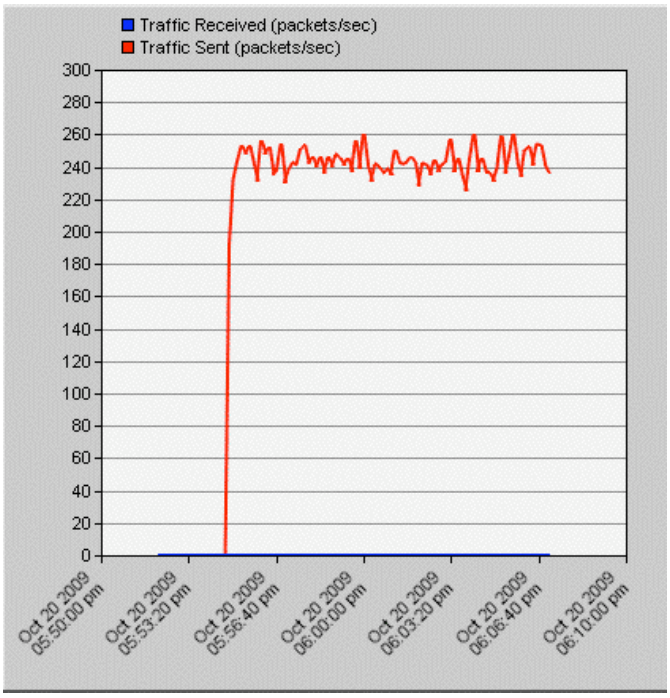

Figure 2.

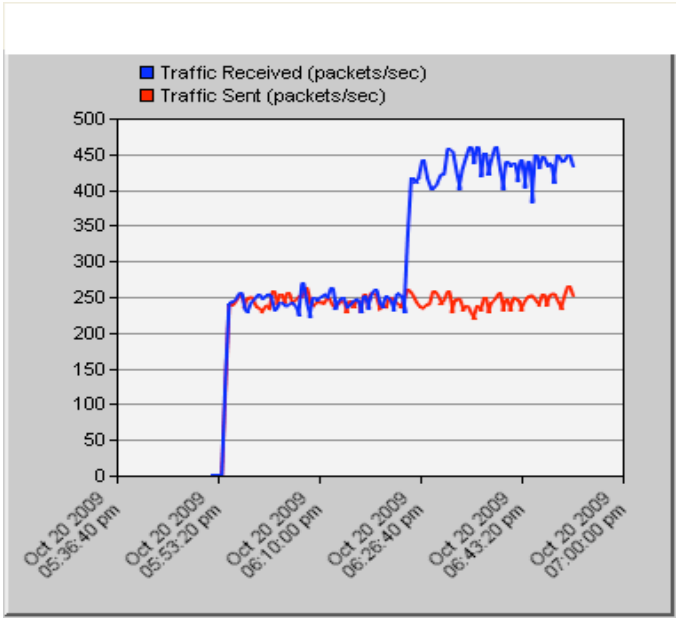

Figure 3.

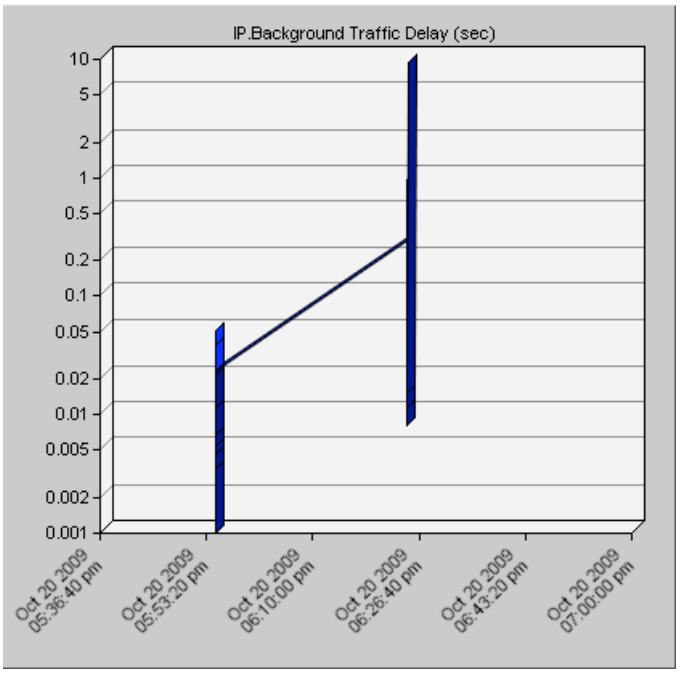

Figure 4.

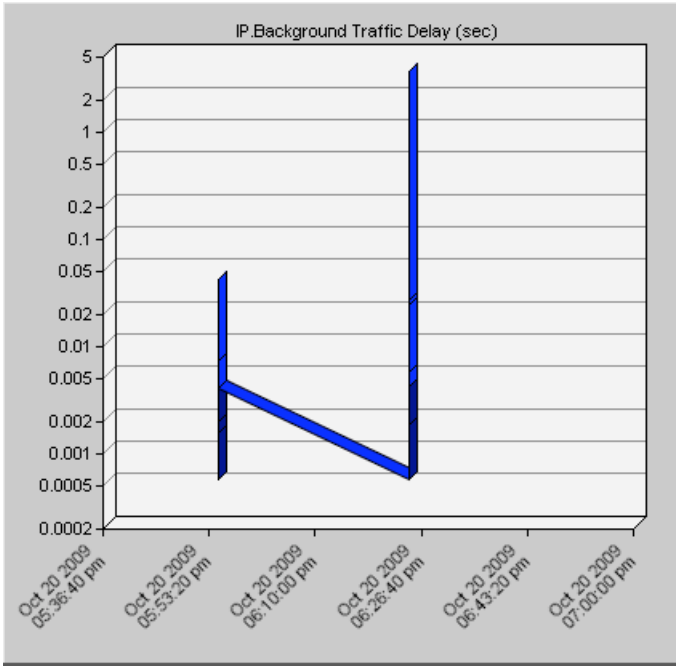

Figure 5.

VI. Conclusion

In our work we demonstrated how AODV control messages are used to severely degrade the performance of nodes in a utilities service territory. We also made recommendations and provided simulations to show that these proposed solutions actually work and can improve network performance (end-to-end) delays during times of increased demand. Looking forward we think that an alternative radio platform like 802.16 with WiMAX connected in a star configuration would eliminate the need for a multihop network. Obviously this eliminates the threat of an attack from within the network that could render several nodes useless. Furthermore utilities with

customers in remote locations could benefit from using WiMAX/802.16 platform by providing multimedia services to their customers. This would help utilities offset loses incurred by implementing a demand response system. For every megawatt a customer sheds the utility will lose a certain percentage of their revenue. Offering data services to customers will indemnify the utility for the losses incurred. To this end we look to simulate the performance of a WiMAX network and determine if such a plan is feasible.

VII. REFERENCES

[1] http://www.getgreenbox.com/
[2] http://blog.mapawatt.com/2009/10/07/list-of-energy-monitoring-tools/
[3] http://news.cnet.com/8301-11128_3-9915135-54.html?tag=mncol;txt
[4] http://www.zigbee.org/
[5] Davis, Mike "SmartGrid Device Security Adventures in a New Medium", *Black Hat USA* 2009
[6] http://earth2tech.com/2009/10/14/ge-backs-solaredge-and-tendril-raises-stake-in-grid-net/
[7] Roosta, Tanya, Shieh, Shiuhpyng, Sastry, Shankar, "Taxonomy of Security Attacks in Sensor Networks and Countermeasures", *The First IEEE International Conference on System Integration and Reliability Improvements*, Hanoi, Vietnam, December 13th-15th 2006
[8] Zheng, Jianliang, Lee J. Myung, Anshel, Michael, "Toward Secure Low Rate Wireless Personal Area Networks," *IEEE Transactions on Mobile Computing*, October 2006 (vol. 5 no. 10)
[9] http://www.opnet.com
[10] Acharya, Arup, Misra, Archan, Bansal Sorav, "A label-switching packet forwarding architecture for multi-hop wireless LANs," *International Workshop on Wireless Mobile Multimedia: Proceedings of the 5th ACM international workshop on Wireless mobile multimedia*, Atlanta GA, 2002
[11] Adibi, Sasan, "MPLS Implementation in Mobile Ad-Hoc Networks (MANETs)," *Ubiquitous Computing and Communication Journal*
[12] RFC 3031
[13] RFC 3561